



Intelligent LAN and Storage Controllers

**Mainframe-class I/O for
PC Servers**

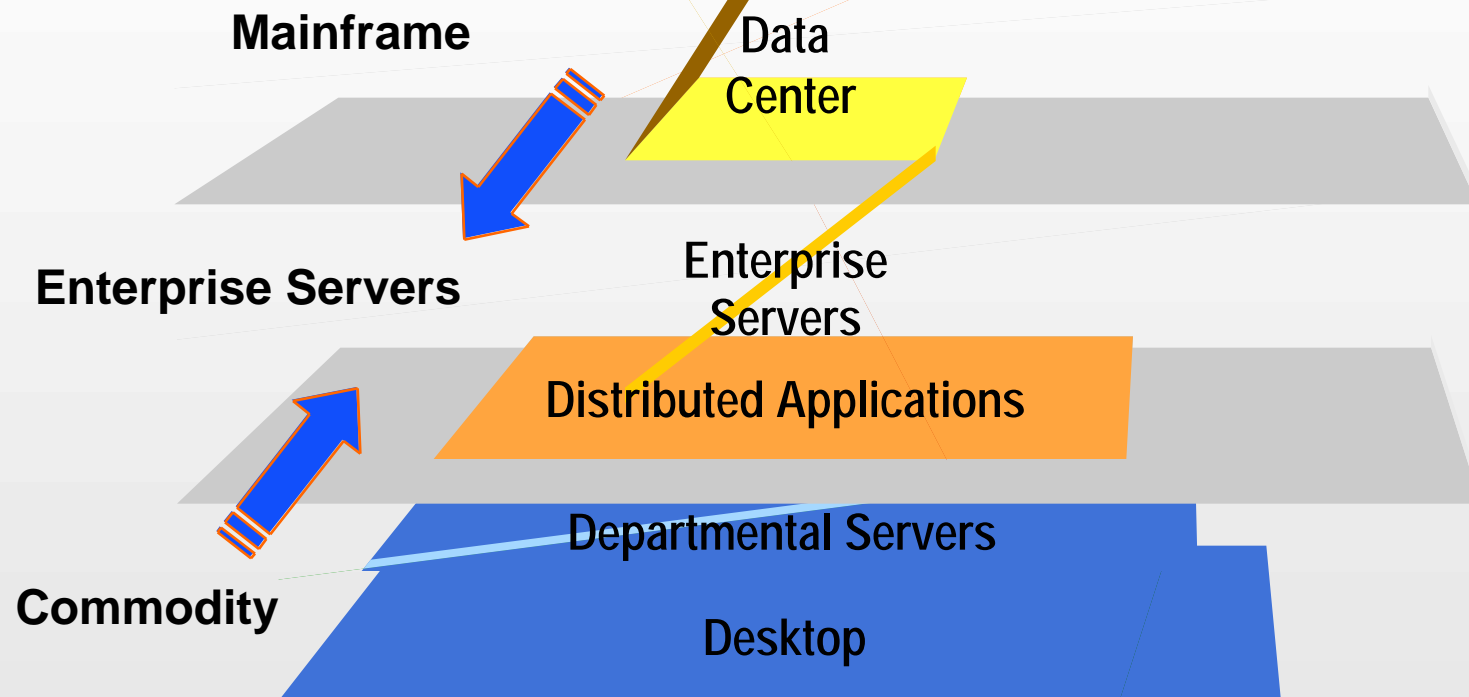
Chris Harrer

Manager, Enterprise Networking

Agenda

- **The Case For Distributing I/O Functionality**
- **An Architectural Model for Distributed TCP/IP**
- **Unisys Networking Products**
- **Performance**
- **Unisys Storage Products**
- **Hardware Platforms**
- **Dynamic RAID Software**
- **Summary**

Converging Technologies

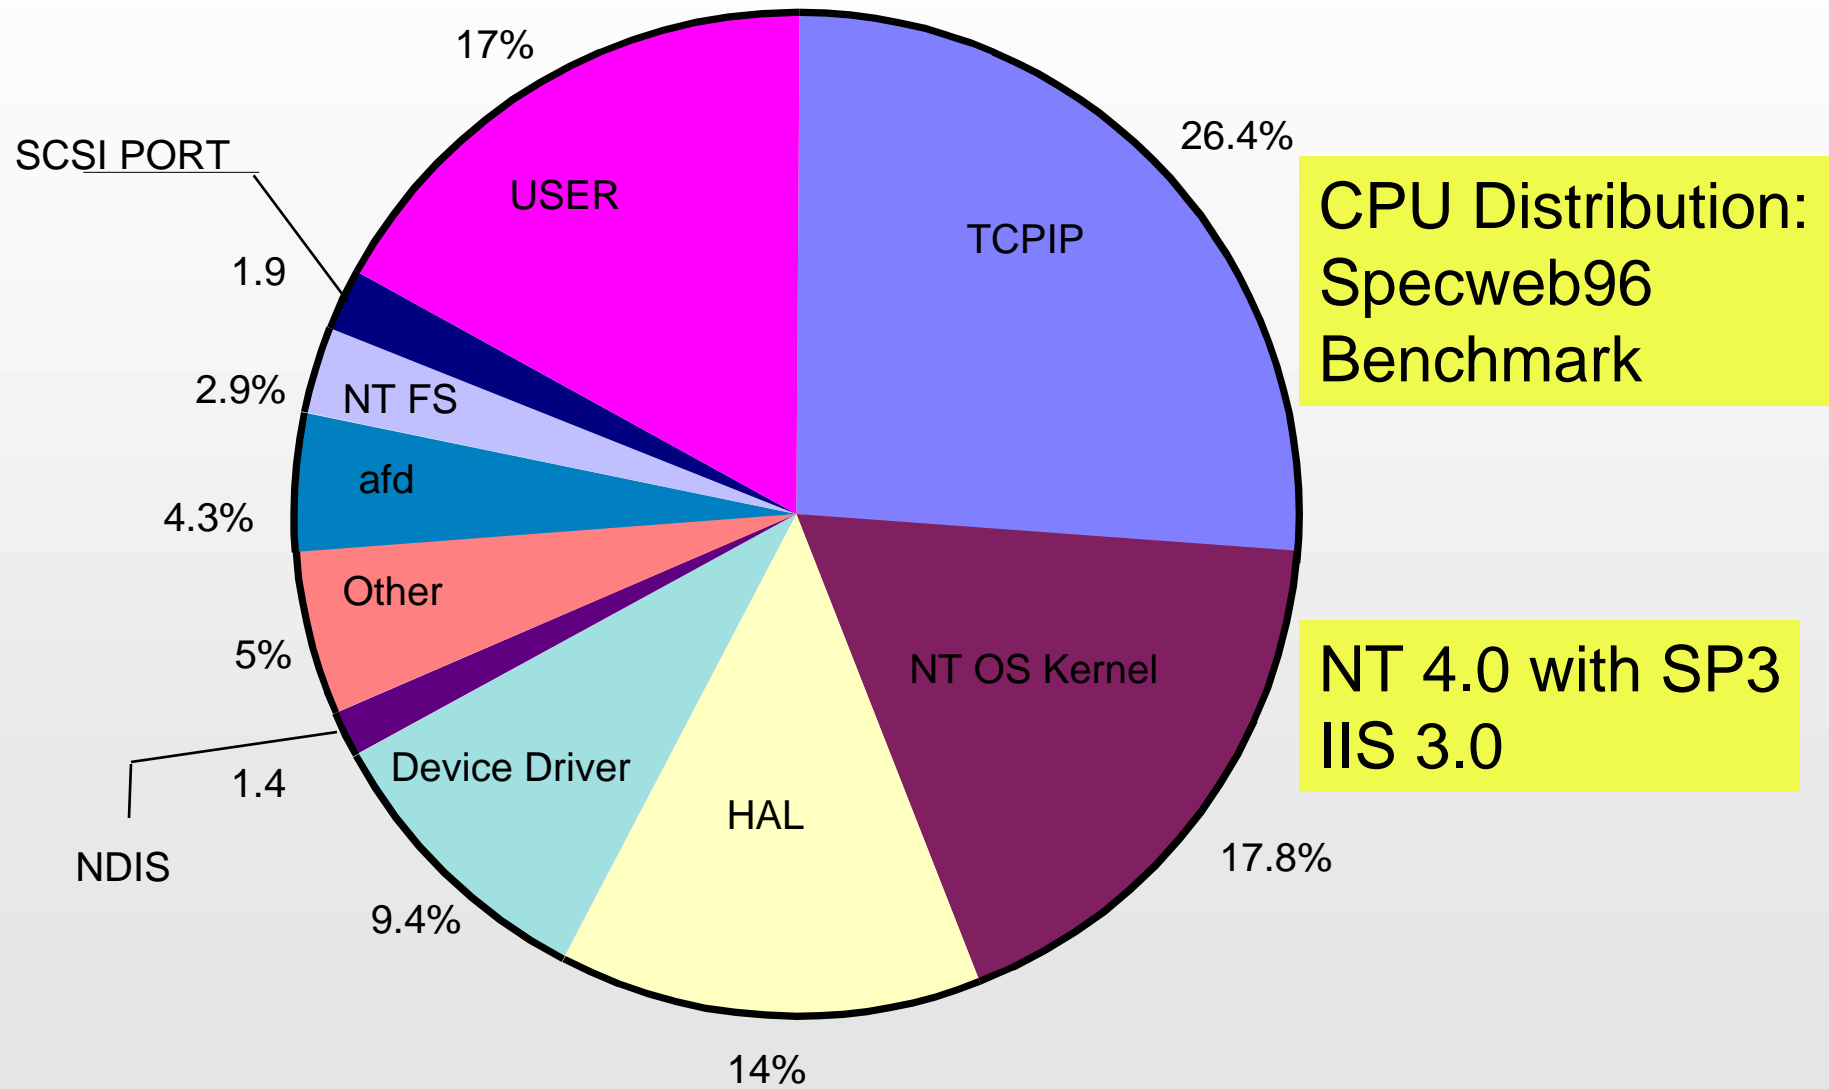


Migrating Distributed I/O to Standard High Volume Servers

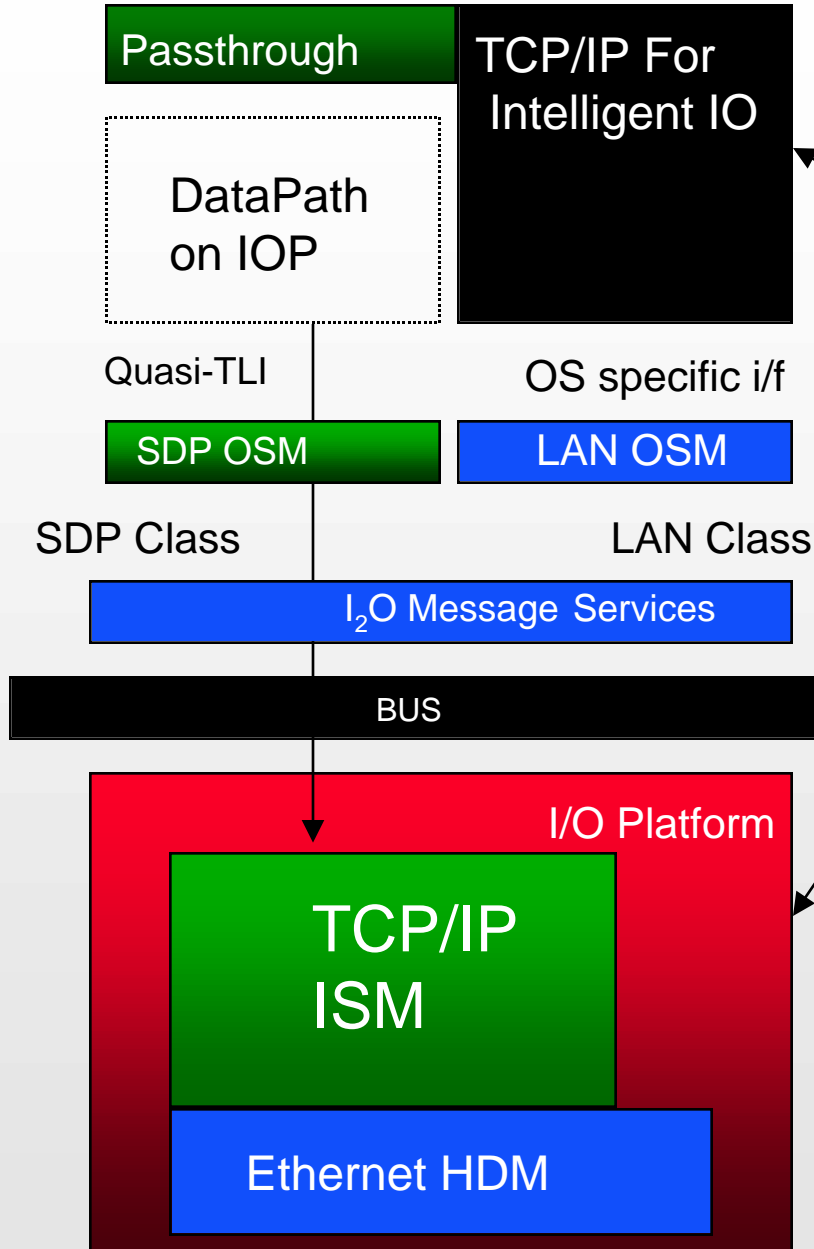
Industry Response

- Market demand for the maturation of Standard High Volume Servers and operating systems for the Enterprise
- Standard High Volume Servers moving to distributed architecture similar to mainframes during the 1980s
- Higher density, lower cost h/w components enabling move to distributed architecture
- Creation of open standard for offloading I/O and networking functionality - I2O

What Functions Can Be Distributed



Architectural Model Distributed TCP/IP Split DataPath



Host TCP/IP stack maintains

- Connection Management
- Route Management
- Application Interface Management
- UDP

Execute data path in IOP

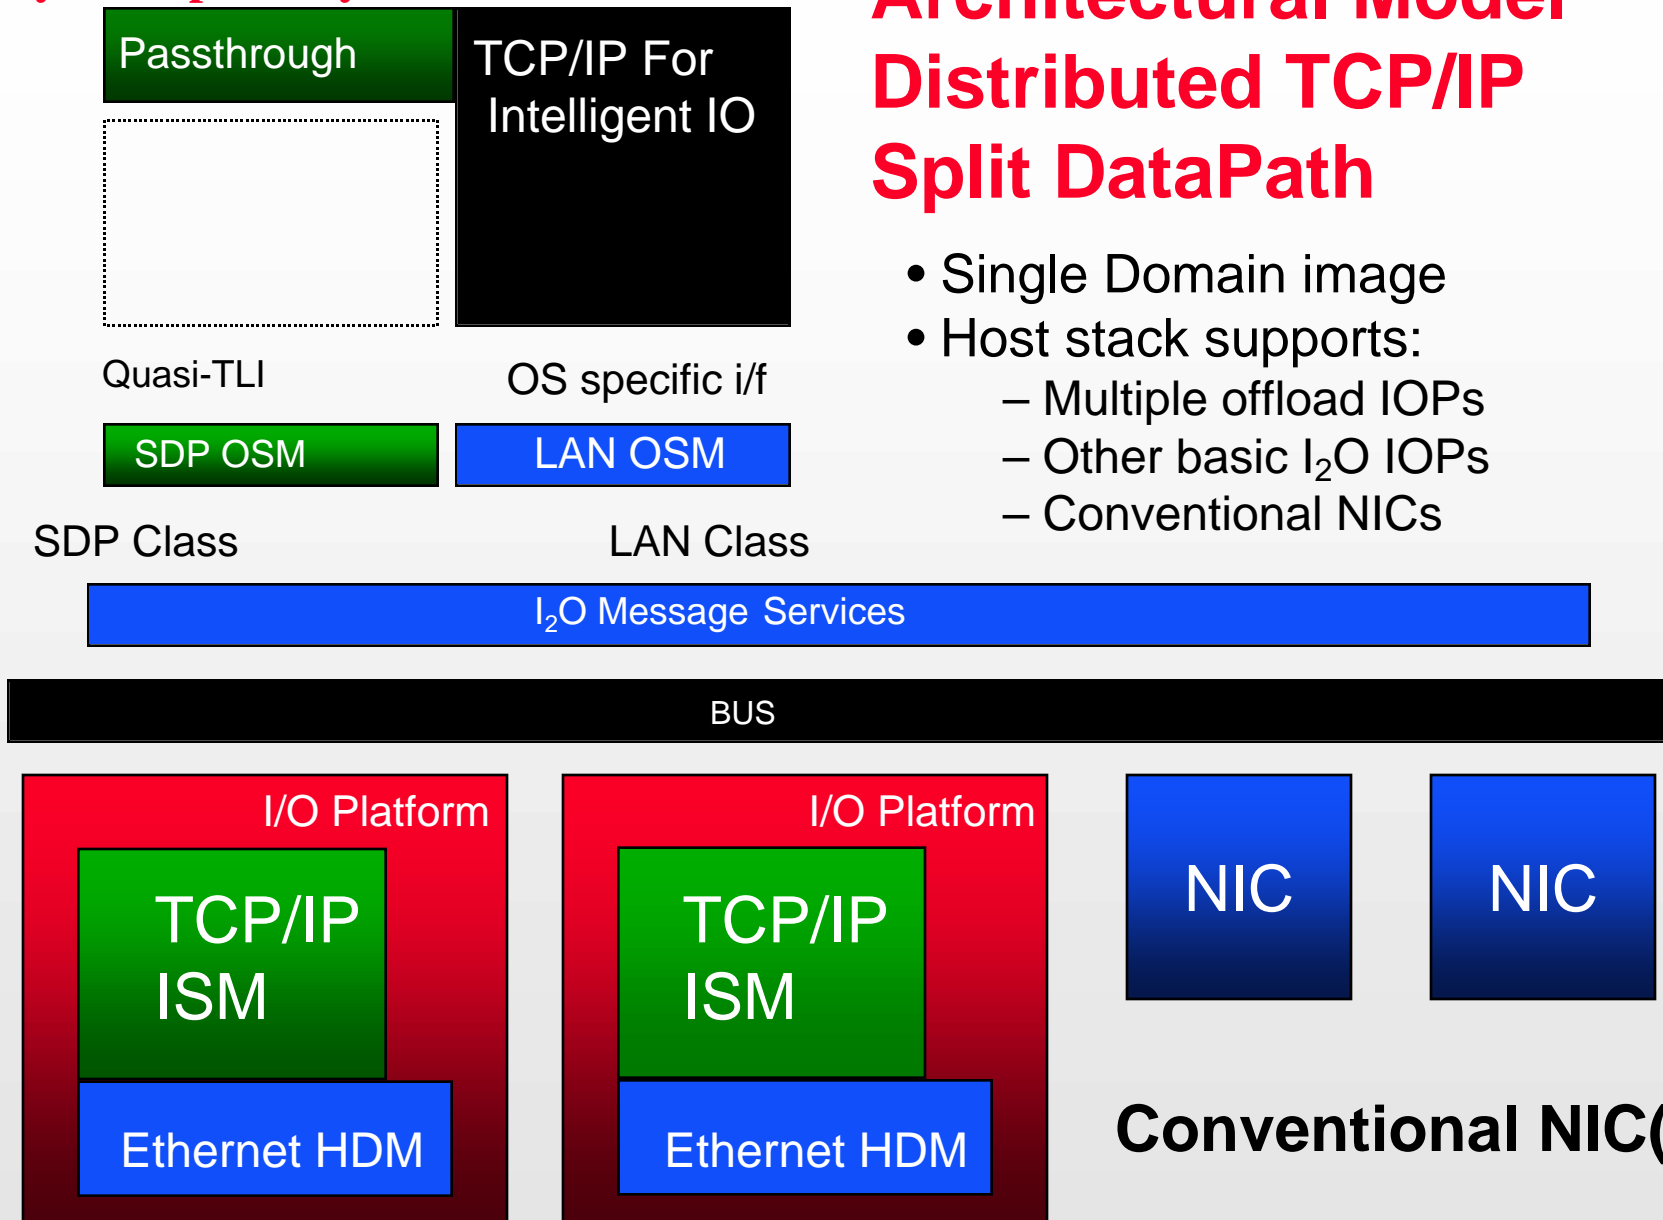
- Checksum
- TCP Segmentation
- Timers
- Retransmissions
- Ack Processing
- Data assembly
- Resequencing
- IP Fragment assembly

Basic I₂O Architect Modules

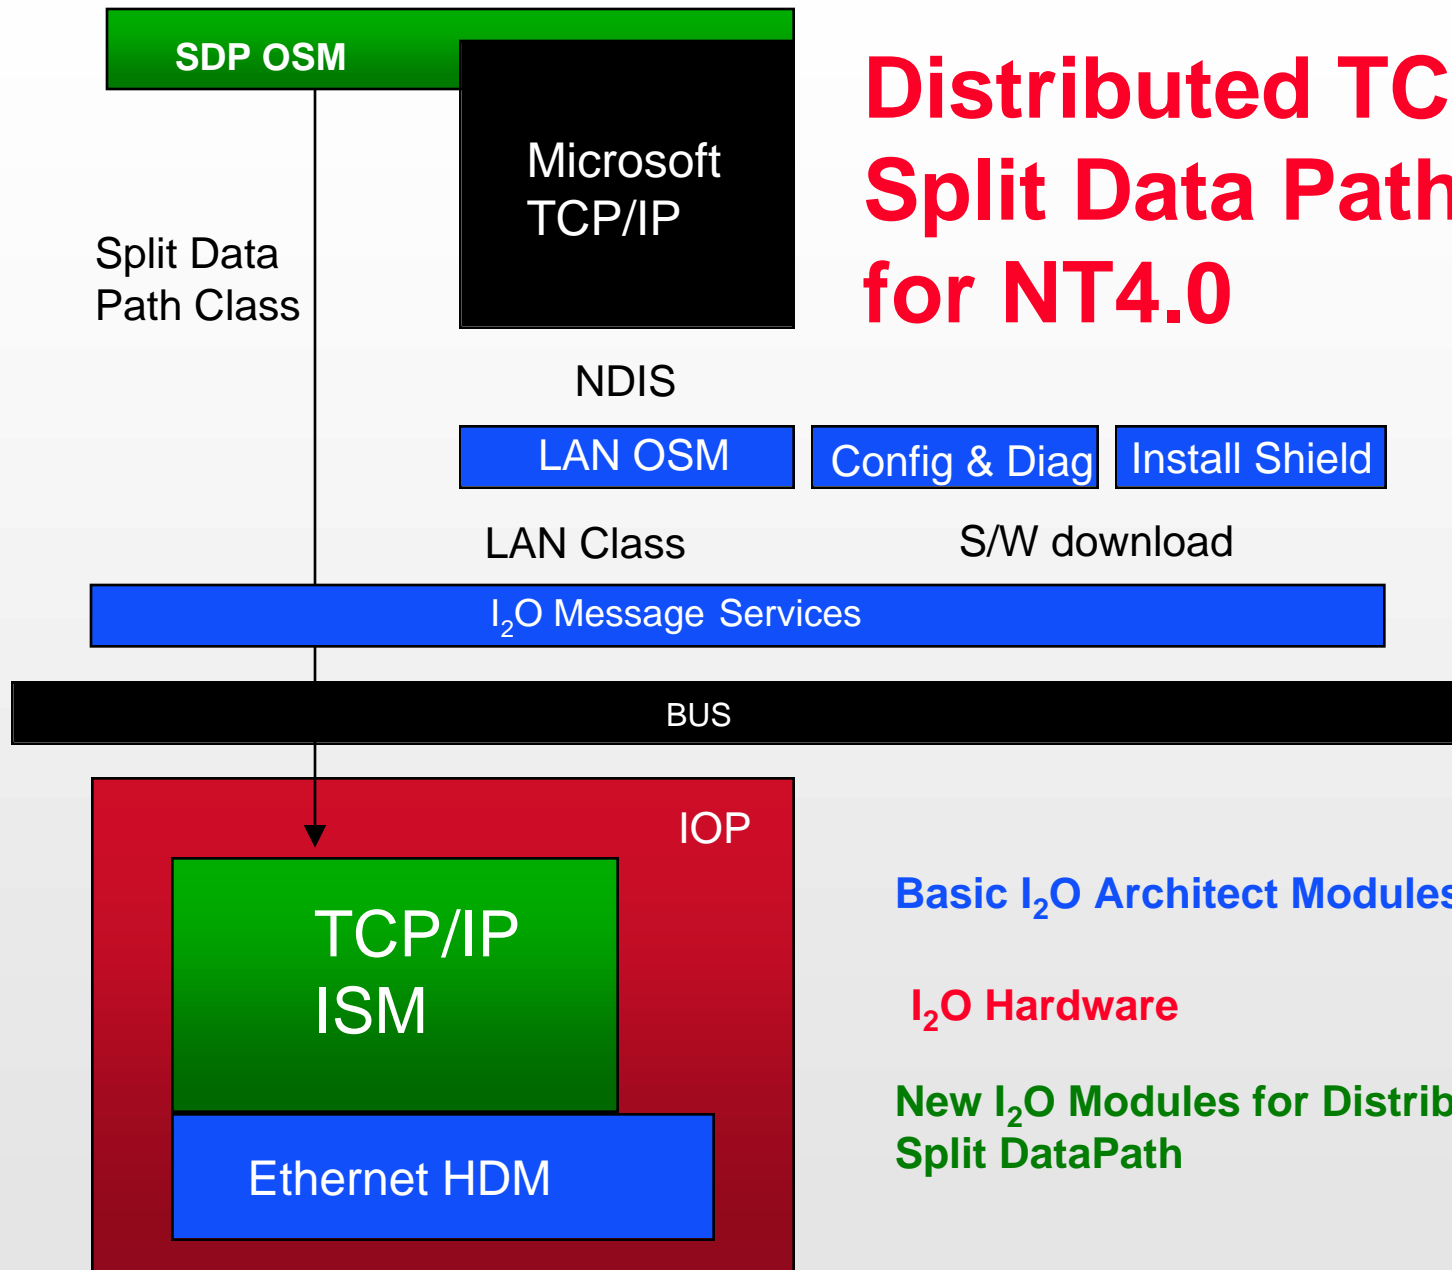
I₂O Hardware

New I₂O Modules for Distributed TCP/IP Split DataPath

Unisys Computer Systems



Distributed TCP/IP Split Data Path for NT4.0



Basic I₂O Architect Modules

I₂O Hardware

New I₂O Modules for Distributed TCP/IP Split DataPath

Split DataPath Software Release 1.0 Features

- NT 4.0 with SP3/SP4
- Basic TCP Data Path Offload (no UDP)
- IPv4
- SNMP MIB II compliance
- Transparent to Applications
- I2O Version 1.5
- Multiple IOPs
- Coexistence w/ conventional NICs
- Multiple devices
- Multiple IP addresses/device

Unisys Computer Systems

IntelliLAN I₂O Fast Ethernet Adapter


Distributed TCP/IP

Extremely Low Host Utilization

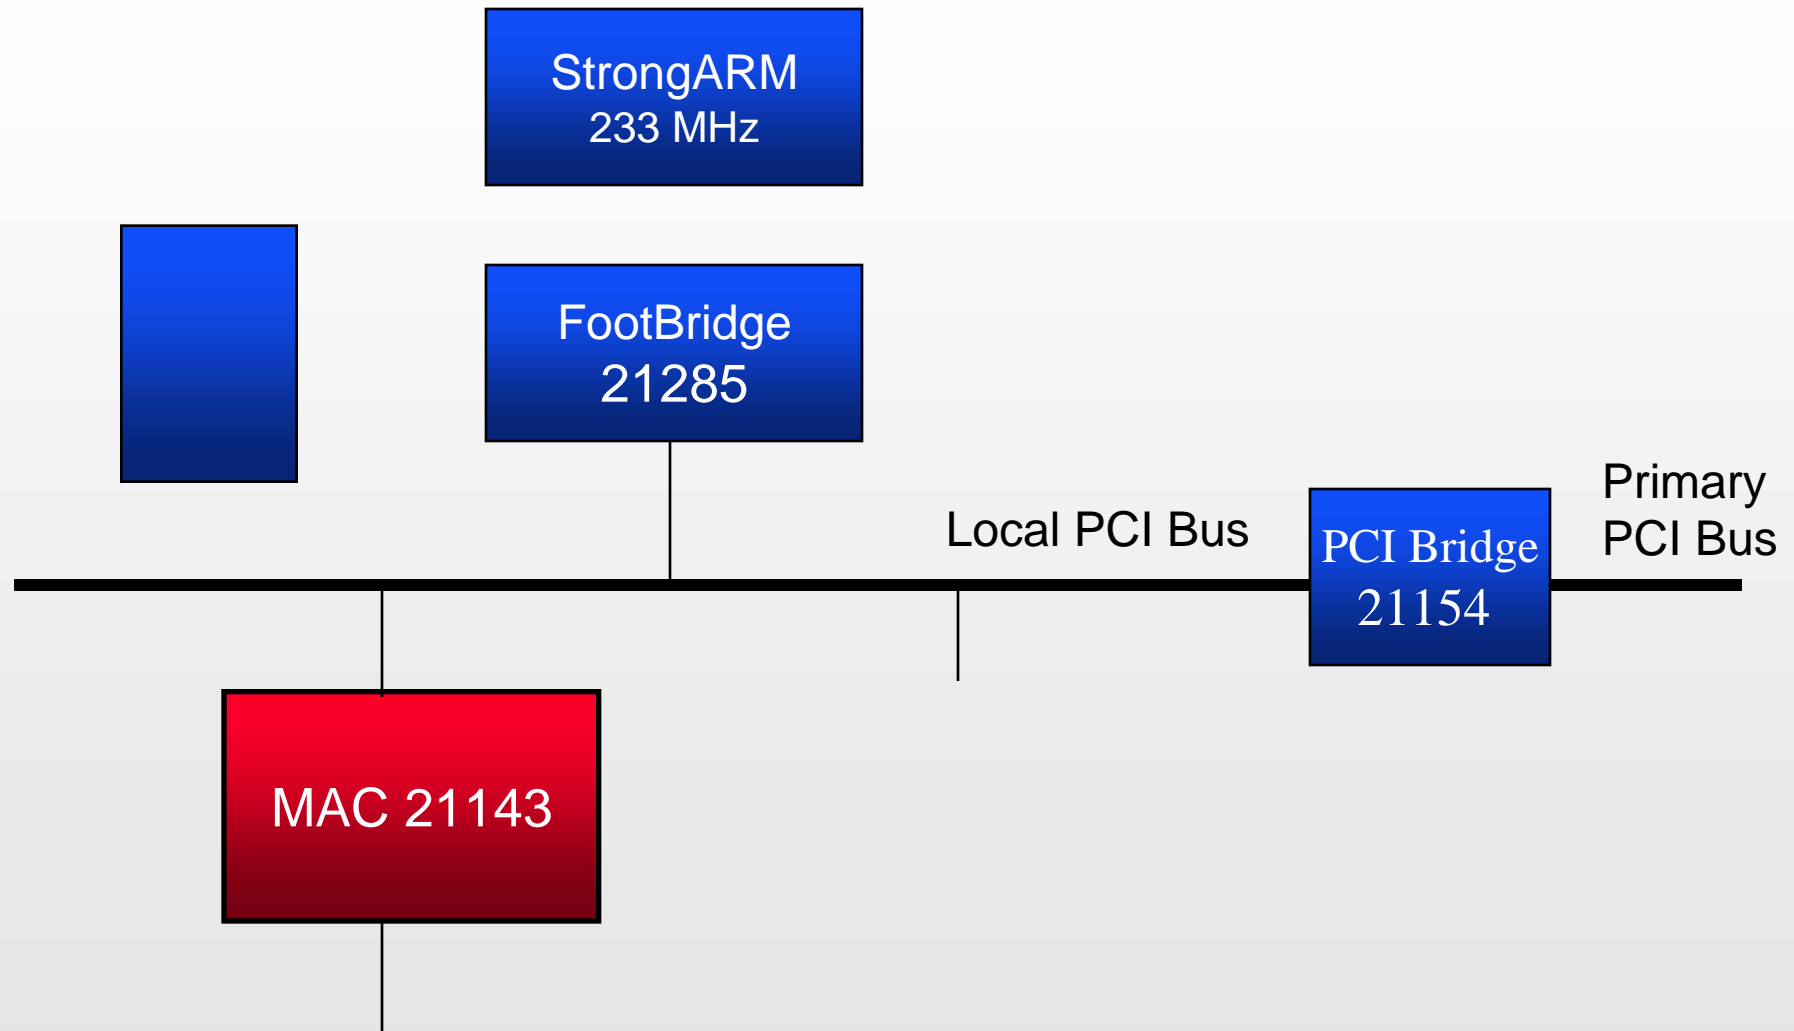
Dual Port

I₂O Based

Wind River IxWorks



Unisys Computer Systems



Performance Objectives

What Are We Measuring

- Server efficiency
 - Host CPU Reduction
- Throughput
 - As opposed to latency
- Compare with industry leading NICs
- Medium to larger configurations
 - As opposed to single connection

Split Data Path Performance Benchmarks

- **Web Bench 2.0**
 - Measures latest HTTP protocol (v1.1) for web servers
 - Multiple GETs per connection
- **SpecWeb96**
 - Measures older HTTP protocol (v1.0) for web servers
 - Still widely used
- **NetPerf**
 - Raw network performance using driver
 - Raw network messaging rate

Split Data Path Performance Benchmarks

- Web Bench 2.0
 - **22 percent less CPU**
- Simulated SpecWeb96 (using Web Bench)
 - **15 percent less CPU**
- **NetPerf**
 - Sending: **46 % less CPU @ 15.0 MBytes/sec**
 - Receiving: **46 % less CPU @ 16.1 MBytes/sec**
 - Send/Receive **45 % less CPU @ 15.1 MBytes/sec**

Split Data Path Performance Summary

- **Paradigms with high data/connection ratios show best performance**
 - Significant improvement in server efficiency
 - Consistent with industry trend as data objects continue to grow
- **Currently measuring high message rate with small payload**
 - Early tests show 20 percent improvement

Performance Directions

- IOP
 - H/W Checksum
 - Improve cache hit rate
 - Larger L1 or small L2 cache
 - Faster memory
 - Alternate parity mechanism allowing optimal bus speeds
- Host
 - Lower management overhead
 - Improved locking
 - Better use of TDI interface and functions
 - Better use of NT services

Post Release 1.0 Candidates

- Improved performance
- NDIS 5.0 offload features
- IPv6
- Recovery from failures to different IOP
- Recovery from IOP Failures
- Jumbo frames
- Port aggregation / load balancing
- Custom timers
- Security



Intelligent Storage Controllers

**Mainframe-class I/O for
Windows NT Servers**

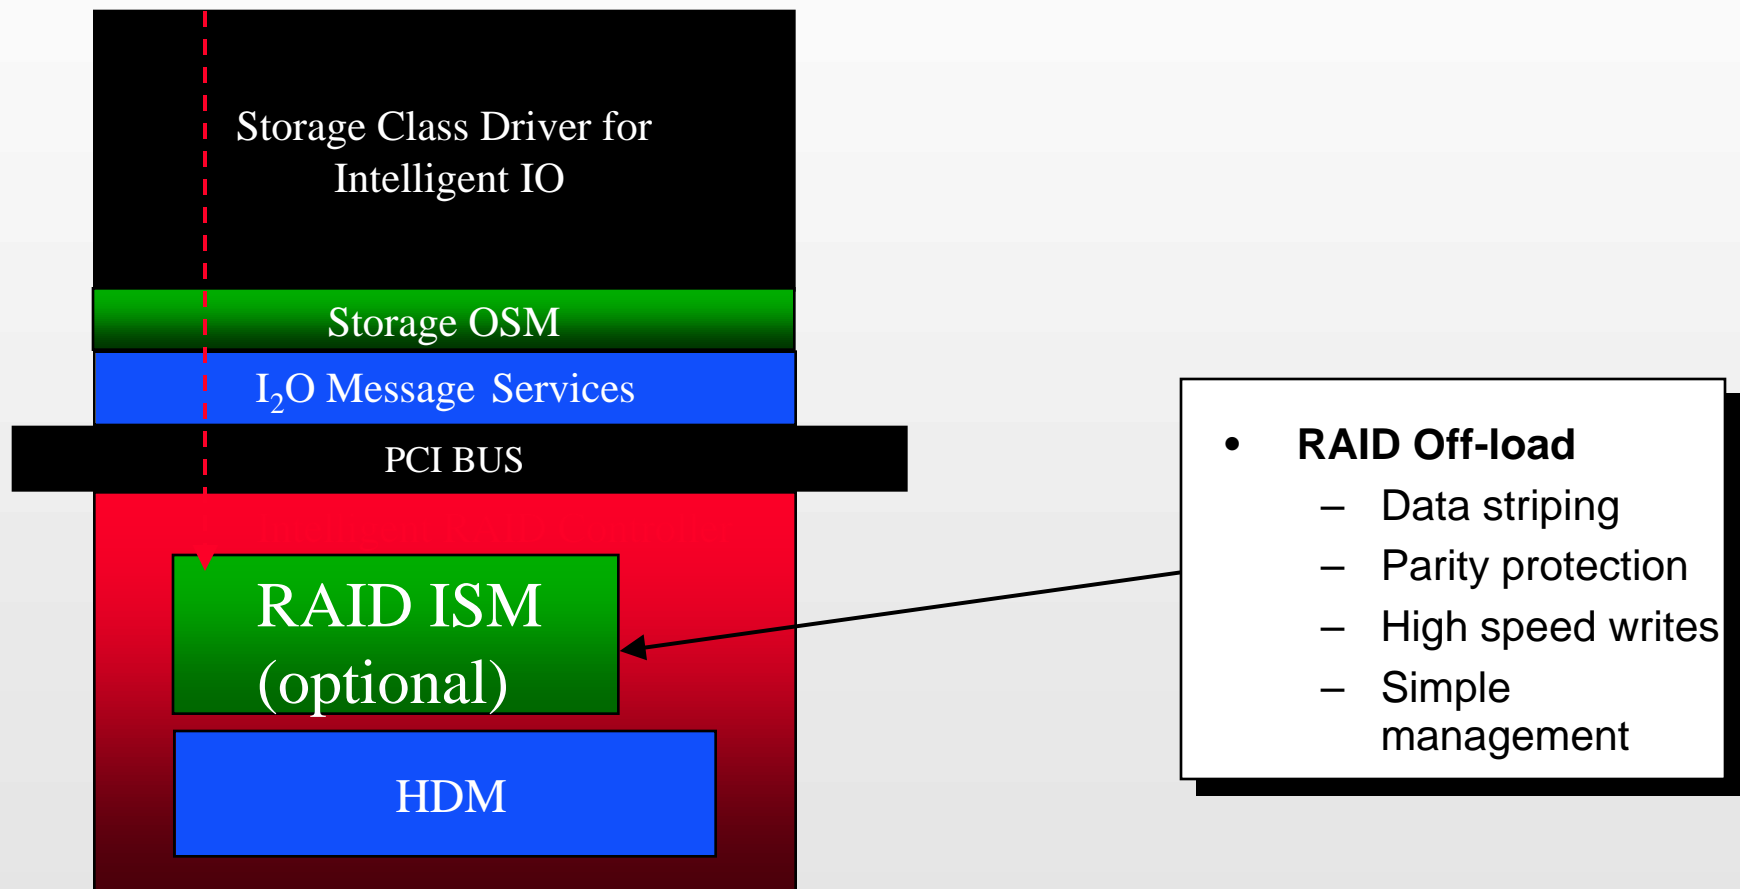
Unisys Intelligent Storage Solutions

- **I₂O Architecture**
- **Windows NT Focus**
- **Intelligent Controller**
 - Fibre Channel
 - Ultra2 Wide SCSI
- **Dynamic RAID (RAID 105)**

Unisys Intelligent Storage Controller Benefits

- **Provides business critical attributes**
- **Simplifies RAID implementation**
- **Supports popular interconnect technologies**
- **Handles greater amounts of information**
- **Adheres to endorsed industry standards**
- **Reduces RAID overhead on system**

Unisys Intelligent Storage Controller Architecture



Unisys Intelligent Storage I₂O Controllers

- **Fibre Channel Controller (3Q99)**
 - dual port
 - 100 MB/sec per loop
 - Supports 125 devices/loop
 - Copper 30 m, Fibre 10 Km
- **Ultra2 Wide SCSI Controller (4Q99)**
 - dual port
 - 80 MB/sec per port

Unisys Intelligent Storage Controller Features

- **NT 4.0 with SP3 or above**
- **PCI form-factor**
- **Dynamic RAID (RAID 105)**
- **Resilient Data Paths (4Q99)**
- **Coexistence with conventional adapters**
- **I₂O Installation, Configuration, Management**

Unisys Computer Systems

IntelliFIBRE Features (non-RAID)

23,000 I/Os Per Second

Extremely Low Host Utilization

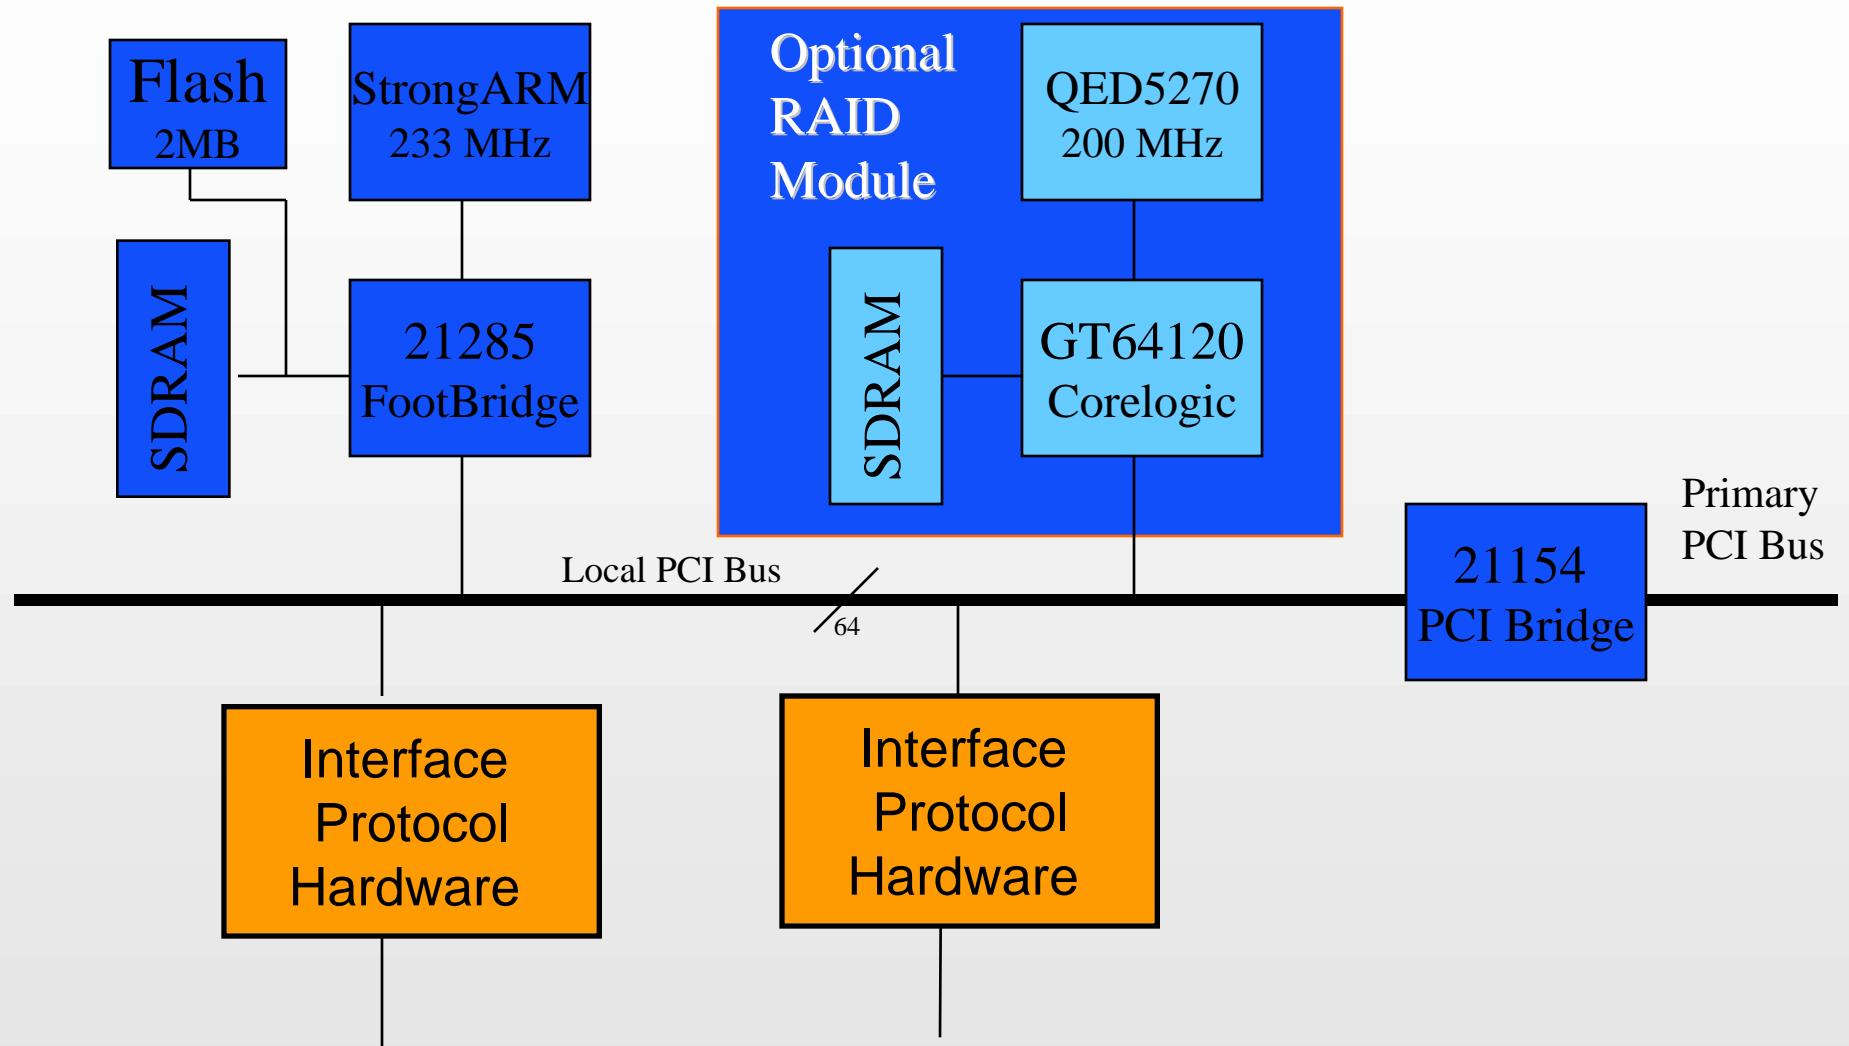
Optional RAID Capability

I₂O Based


Wind River IxWorks



Storage Controller Hardware Structure




Software - Dynamic RAID




- RAID 5 storage efficiency with RAID 10 performance!
- **14,000 I/Os Per Second**
- Add protected capacity in less than a minute
- Add capacity 1 or more drives at a time
- Simple and flexible
- Supports Hot or Stand-By spares

RAID 10 Overview




- RAID 10 is based on the concept of mirroring. Using two drives, data is written to the originally targeted drive and a copy of the data is written to its mirrored pair.
- When data is requested by a host, it can be retrieved from either side of the mirror.
- High performance
- Storage capacity requirement 2x data requirement

RAID 5 Overview



- RAID 5 stripes data and parity across all the drives in an array.
- For example, the drives on the left form an array. The white rings equal a single stripe. The black, yellow and red are individual stripes as well.
- Poorer write performance.
- Storage capacity requirement 1.25x data requirement

Dynamic RAID Protection Overview



For example, pictured above, based on 9GB drives, capacity is as follows:

RAID 10 Allocation

$$9\text{GB} - 10\% = 8.1\text{GB}$$

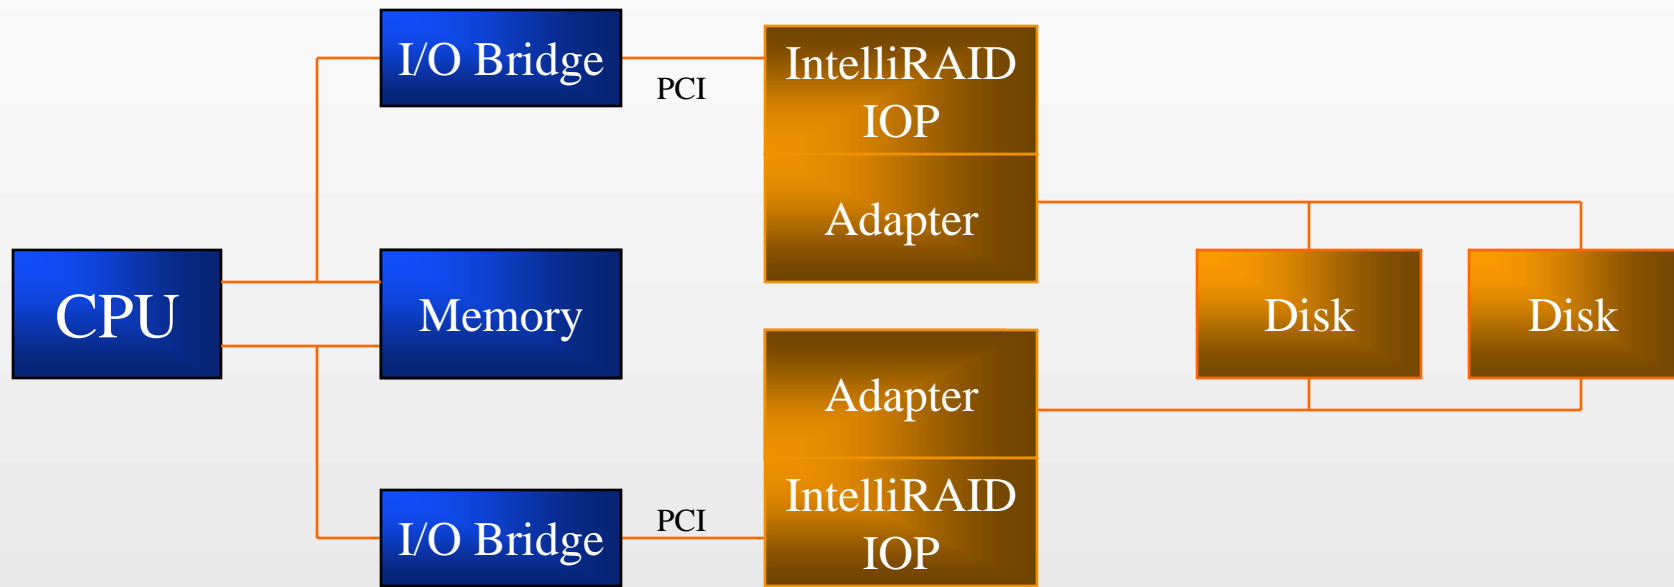
RAID 5 (N-1)

$$32.4\text{GB} (8.1 \times 4) - 8.1 (8.1 \times 1) = 24.3\text{GB}$$

useable capacity

- From each drive, 10% is taken for the RAID 10 Block Pool.
- The remaining capacity of each of the drives in the array is based on one drive for parity and the remaining drives for data; N-1.
- Performance approaches RAID 10
- TCO approaches RAID 5
- Unisys has filed for three patents

Resilient Data Paths



Summary

- **Migrate Unisys heritage of resilient and scalable I/O to PC Servers**
- **Increase system efficiency**
- **Reduce overall cost**